


Firewalls



Introduction

Firewalls are a key part of keeping networked computers safe and secure. All computers deserve the protection of a firewall, whether it's the thousands of servers and desktops that compose the network of a Fortune 500 company, a traveling salesperson's laptop connecting to the wireless network of a coffee shop, or your grandmother's new PC with a dial-up connection to the Internet. This article covers the design, deployment, and use of both network and host-based firewalls (also called personal firewalls). Although home users have traditionally used only host-based firewalls, recent trends in security exploits highlight the importance of using both types of firewalls together. Traditional firewall architectures protect only the perimeter of a network. However, once an attacker penetrates that perimeter, internal systems are completely unprotected. Hybrid worms, in particular, have penetrated corporate networks through email systems, and then have spread quickly to unprotected internal systems. Applying host-based firewalls to all systems, including those behind the corporate firewall, should now be standard practice.

The Nature of Today's Attackers

Who are these "hackers" who are trying to break into your computer? Most people imagine someone at a keyboard late at night, guessing passwords to steal confidential data from a computer system. This type of attack does happen, but it makes up a very small portion of the total network attacks that occur. Today, worms and viruses initiate the vast majority of attacks. Worms and viruses generally find their targets randomly. As a result, even organizations with

little or no confidential information need firewalls to protect their networks from these automated attackers.

If a worm or a virus does find a security vulnerability and compromises your system, it can do one of several things. To begin with, it will almost always start looking for other systems to attack so that it can spread itself further. In this case, you become one of the bad guys—because the worm or virus is using your computer to attack other systems on your internal network and the Internet, wasting your computing resources and bandwidth. Even though the worm or virus won't know what to do with your confidential data, chances are good that it will open a new back door into your system to allow someone else to further abuse your computer and compromise your privacy. Worms and viruses have dramatically increased the need for network security of all kinds—especially the need for host-based firewalls.

Individuals still launch some attacks, though, and these are generally the most dangerous. The least worrisome attackers focus on crashing computers and networks by using Denial of Service (DoS) attacks. Others might be looking for confidential data that they can abuse for profit, such as sales contacts, financial data, or customer account information. Still others might be amassing hundreds or thousands of computers from which to launch a distributed attack against a single network on the Internet.

The Firewall to the Rescue

In the physical world, businesses rely on several layers of security. First, they rely on their country's government and military forces to keep order. Then, they trust their local police to patrol the streets and respond to any crimes that occur. They further supplement these public security mechanisms by using locks on doors and windows, employee badges, and security systems. If all these defenses fail and a business is a victim of a crime, the business's insurance agency absorbs part of the impact by compensating the business for a portion of the loss.

Unfortunately, the state of networking today lacks these multiple levels of protection. Federal and local governments do what they can to slow network crime, but they're far from 100 percent effective. Beyond prevention, law enforcement generally only responds to the most serious network intrusions. The

average Internet-connected home or business is attacked dozens of times per day, and no police force is equipped to handle that volume of complaints. Losses from computer crime are hard to quantify and predict, and as a result most business insurance policies do little to compensate for the losses that result from a successful attack.

The one aspect of physical security, however, that isn't missing from network security is the equivalent of door locks, employee badges, and security systems: firewalls. Just as you lock your car and home, you need to protect your computers and networks. Firewalls are these locks, and just like in the physical world, they come in different shapes and sizes to suit different needs. The famous Jargon Dictionary has a great definition for *firewall*: "a dedicated gateway machine with special security precautions on it, used to service outside network connections and dial-in lines." Firewalls serve two useful purposes: they filter what traffic comes into your network from the outside world, and they control what computers on your network may send there.

It's important to understand one thing, however. No firewall—whether a small, free host-based firewall or a multiple-thousand-dollar enterprise firewall array—will make your computers impervious to attack. Firewalls, like locks and walls and moats and dragons, create barriers to attack—they get in the way of someone trying to take control. By making it difficult for attackers to get into your computer, by making them invest lots of time, you become less attractive. Firewalls very effectively block most bad guys from compromising an individual computer. But it's impossible to fully prevent every intrusion: All software has bugs, and someone might find an obscure bug in your firewall that allows them to pass through. In a nutshell, there's no such thing as absolute security. How much you invest in firewalls should be a function of how much you have to lose if an attack is successful.

Types of Firewalls

There are two main types of firewalls: network firewalls and host-based firewalls. Network firewalls, such as the software-based Microsoft's Internet Security and Acceleration (ISA) Server or the hardware-based Nortel Networks Alteon Switched Firewall System, protect the perimeter of a network by watching traffic that enters and leaves. Host-based firewalls, such as Internet Connection

Firewall (ICF—included with Windows XP and Windows Server 2003), protect an individual computer regardless of the network it's connected to. You might need one or the other—but most businesses require a combination of both to meet their security requirements.

How a Firewall Works

The sections that follow provide background information about network traffic and how firewalls filter traffic. This information applies to all types of firewalls.

Basic TCP/IP Flow

This section describes how TCP/IP packages its information, to show how firewalls decide to allow or deny traffic. TCP/IP traffic is broken into packets, and firewalls must examine each packet to determine whether to drop it or forward it to the destination. Figure 1 shows a simplified breakdown of a packet with the following three key sections: the IP header, the TCP or UDP header, and the actual contents of the packet. The IP header contains the IP addresses of the source, which is the sender, and the destination, which is the receiver. The TCP or UDP header contains the source port of the sender and the destination port of the receiver to identify the applications that are sending and receiving the traffic. In addition, TCP headers contain additional information such as sequence numbers, acknowledgment numbers, and the conversation state. The destination TCP or UDP ports define the locations for delivery of the data on the server when the packet reaches its destination.

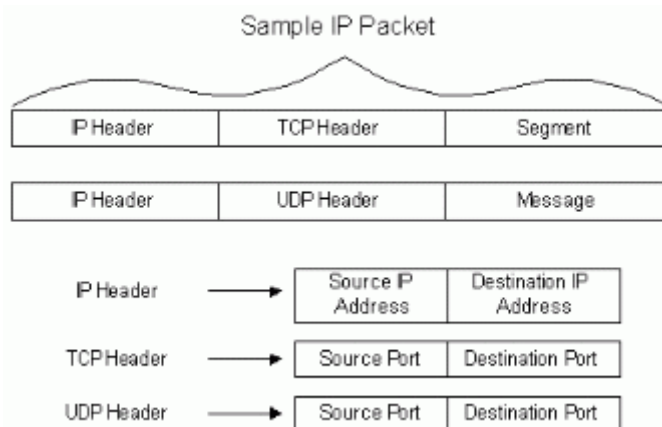


Figure 1: An IP packet contains a header useful to firewalls.

[See full-sized image.](#)

It's important to appreciate the communication flow of a TCP/IP conversation when configuring the firewall. When a browser, for example, sends an HTTP request to a Web server, the request contains the identity of the client computer, the source IP address, and the source port that the request went out on. The source port of the client identifies the client application that sent the request—in this case, the browser. When the Web server sends a response, it uses the client's source port as the destination port in the response. The client operating system recognizes the port number as belonging to a session the browser application started, and gives the data to the browser. The source port for a client is typically a value greater than 1024 and less than 5000.

Packet Filtering

The primary purpose of a firewall is to filter traffic. Firewalls inspect packets as they pass through, and based on the criteria that the administrator has defined, the firewall allows or denies each packet.

Firewalls block everything that you haven't specifically allowed. Routers with filtering capabilities are a simplified example of a firewall. Administrators often configure them to allow all outbound connections from the internal network, but to block all incoming traffic. So, a user on the internal network would be able to download email without a problem, but an administrator would need to customize the router configuration to connect to your home PC from work by using Remote Desktop. Other applications that might require special firewall configuration are WebCam servers, collaboration software, and multiplayer online games.

You use packet filters to instruct a firewall to drop traffic that meets certain criteria. For example, you could create a filter that would drop all ping requests. You can also configure filters with more complex exceptions to a rule. For example, a filter might assist with troubleshooting the firewall by allowing the firewall to respond to ping requests coming from a monitoring station's IP address. By default, Microsoft ISA Server doesn't respond to ping queries on its external interface. You would need to create a packet filter on the ISA Server computer for it to respond to a ping request.

The following are the main TCP/IP attributes used in implementing filtering rules:

- Source IP addresses
- Destination IP addresses

- IP protocol
- Source TCP and UDP ports
- Destination TCP and UDP ports
- The interface where the packet arrives
- The interface where the packet is destined

If you've configured the firewall to allow all traffic by default, you can use filters to block specific traffic. If you've configured the firewall to deny all traffic, filters allow only specific traffic through. A common packet-filtering configuration is to allow inbound DNS requests from the public Internet so that a DNS service can respond.

Developers have designed most applications to work properly with both routers and host-based firewalls, but some might require you to configure your firewall to allow the application to communicate. Fortunately, firewalls are very common, and any application that requires a firewall should include information about how to configure your firewall. Host-based firewalls are easier to configure than network firewalls and generally include a wizard to walk you through the configuration process. Many host-based firewalls automatically prompt you the first time any application attempts to use the Internet—whether the connection is inbound or outbound. While using a host-based firewall, you might even notice applications that you didn't know accessed the Internet, such as Microsoft Word. Figure 2 shows the filter configuration screen for ICF:

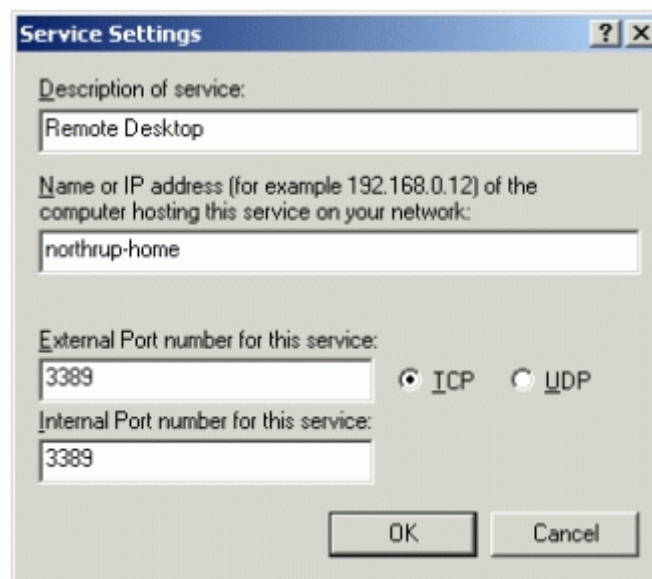


Figure 2: ICF allows custom filters to be created.

Figure 3 shows the filter configuration screen for a third-party firewall application, ZoneAlarm Pro:

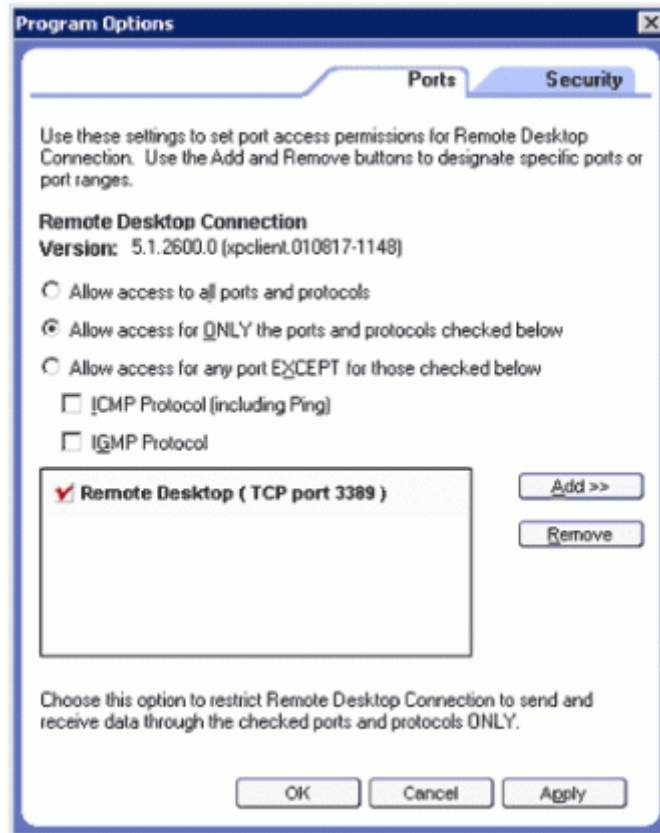


Figure 3: ZoneAlarm also allows custom filters to be created.

Both of these examples demonstrate enabling the Remote Desktop feature in Windows XP, which uses TCP port 3389. You won't be able to connect unless you open up port 3389 to permit the Remote Desktop capability to pass through. For more information, see the Microsoft Knowledge Base article 308127, How to Manually Open Ports in Internet Connection Firewall in Windows XP.

Most modern firewalls are friendly enough so that they hide the port numbers from you. For example, ICF allows you to choose the names of the applications that you want to allow through. However, it's very common to need to add an application to the list of allowed traffic. To add an application, you need to know the port number that the application uses.

Table 1 shows a list of port numbers for commonly used applications. As mentioned earlier, ports can be either a TCP port or a UDP port. Most applications use TCP ports. However, DNS uses UDP, and without DNS, you wouldn't be able to find Web sites on the Internet.

Table 1 Common Port Numbers

Service	Port
Web server	80/tcp
SSL (Secure Sockets Layer) Web server	443/tcp
FTP	21/tcp
POP3	110/tcp
SMTP	25/tcp
Remote Desktop (Terminal Services)	3389/tcp
IMAP3	220/tcp
IMAP4	143/tcp
Telnet	23/tcp
SQL Server	1433/tcp
LDAP	389/tcp
MSN Messenger	1863/tcp
Yahoo! Messenger	5050/tcp
AOL Instant Messenger and ICQ	5190/tcp
IRC (Internet Relay Chat)	6665-6669/tcp
DNS	53/udp

To use TCP/UDP port-filtering tools effectively, configure the filtering tool to accept requests through each port that your server applications require, and to refuse requests from all other TCP or UDP ports. Making a careful determination

of your applications' TCP/UDP port requirements and setting your filtering tools accordingly allows you to avoid mistakes that would deny access to the services you're trying to provide. Filtering out all traffic to other TCP and UDP ports eliminates unnecessary exposure to attack.

Filtering Based on Source and Destination

Some types of firewalls can filter traffic based on source or destination IP address. IP addresses are the telephone numbers of the Internet: They're the unique, numeric label that identifies a single host's location. Filtering based on source or destination address is useful because it enables you to allow or deny traffic based on the computers or networks that are sending or receiving the traffic.

This is useful in two ways. First, you can configure firewalls to block specific Web sites. Blocking Web sites by name is a form of destination filtering. Second, firewalls can allow or deny traffic based on the computer sending the request. This allows administrators to disable instant messaging from the computer in one organization, while allowing the same protocol from a different set of computers. Source filtering also allows you to give greater access to users on internal networks than those on external networks. It's common to use a firewall to block all requests sent to an internal email server except those requests from users on the internal network. You can also use source filtering to block all requests from a specific address—for example, to block traffic from an IP address identified as having attacked the network.

Stateful Inspection Filtering

Stateful inspection is the process of inspecting packets as they reach the firewall and maintaining the state of the connection by allowing or disallowing packets to pass based on the access policy. To further help you understand how state is maintained, Figure 4 shows how a conversation between a client and a server takes place through the ISA Server computer. In this scenario, Web Publishing has been configured on the ISA Server computer to support redirecting external Internet requests on port 80 to the internal IIS server:

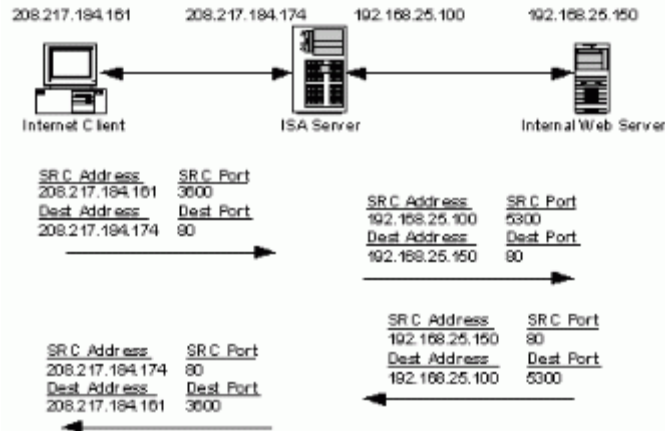


Figure 4: Sample conversation through ISA Server

This is the flow of the conversation:

1. The Internet client initiates an HTTP request to the Web server and sends an IP packet with the source and destination address and ports.
2. The ISA Server computer receives the request for the Web server.
3. ISA Server then modifies the packet, replacing the source address and port with its own internal address, and changes the destination IP address to the address of the real IIS server.
4. ISA Server adds the source and destination ports and addresses into its own table to keep track of the conversation.
5. ISA Server sends the modified packet to the internal IIS server.
6. The IIS server responds to the request by using ISA Server as the destination address and TCP port 5300.
7. ISA Server receives the packet from the IIS server and looks in its table for 5300, which maps to the Internet client.
8. ISA Server then modifies the packet and replaces the IIS server's source IP address and port with its own source IP address and port.
9. ISA Server then changes the destination IP address and TCP port to that of the Internet client.
10. The Internet client listens for a response on TCP port 5100.

In addition to maintaining the TCP or UDP conversation based on IP addresses and ports, ISA Server also checks the TCP flags, the sequence and acknowledgment numbers within the TCP header fields for TCP conversations. The flags represent the state of the conversation, whether it's the beginning of a

conversation (SYN), the middle of a conversation (ACK), or the end of the conversation (FIN). If any of the flags are out of sequence, ISA Server blocks the connection. The sequence and acknowledgment fields provide the information to ensure that the next packet received in the conversation is the correct one. Once again, any request that doesn't fit the state of the conversation is blocked.

Application-Layer Filtering

Application-layer firewalls can understand the traffic flowing through them and allow or deny traffic based on the content. Host-based firewalls designed to block objectionable Web content based on keywords contained in the Web pages are a form of application-layer firewall. You also use application-layer firewalls to inspect packets bound for an internal Web server to ensure the request isn't really an attack in disguise.

Currently, the ability to inspect a packet's contents is one of the best ways to distinguish between firewall products. ICF lacks this feature. However, most business-oriented firewalls do include this capability.

ISA Server is also an application-level proxy that's able to read data within packets for a particular application and perform an action based on a rule set. In addition, ISA Server comes with predefined application filters that inspect each packet and block, redirect, or modify the data within the packet. For instance, you can implement Web-routing rules that tell the ISA Server computer to redirect an HTTP request to a certain internal IIS server, based on the URL in the packet. Another example is the DNS intrusion-detection filter. This filter blocks packets that aren't valid DNS requests, or that fit common types of DNS attacks. You can invoke application filtering on ISA Server when Web Publishing or Server Publishing is configured.

Logging

Firewalls don't prevent attacks; they simply reduce the likelihood of a break-in. When you deploy a firewall, you'll still get just as many attacks as you always did—you just won't have to worry about them as much. All firewalls provide some capability for logging these attacks for later, manual review. This allows administrators to watch for attacks that are out-of-the-ordinary. It's also useful for forensics purposes. If an attacker does manage to defeat your firewall, you can

refer to the firewall's log and gather information to determine how the attacker carried out the attack. This log can be useful to law enforcement officials, if they're involved in a related investigation.

Intrusion Detection

Intrusion detection is an advanced firewall feature, and many firewalls (such as ICF) lack this feature. Intrusion detection systems (IDSs) can identify attack signatures or patterns, generate alarms to alert the operations staff, and cause the routers to terminate the connection with the hostile sources. These systems can also prevent DoS attacks. A DoS attack occurs when a user sends fragments of TCP requests, masked as legitimate TCP requests, or sends requests from a bad IP source. The server can't handle so many requests and displays a DoS message to legitimate site users. IDSs provide real-time monitoring of network traffic and implement the "prevent, detect, and react" approach to security.

Although IDSs are necessary to meet security requirements for many businesses and some home users, their use has downsides that you should take into account:

- IDSs are processing-intensive and can affect the performance of your site.
- IDSs are expensive.
- IDSs can sometimes mistake normal network traffic for a hostile attack and cause unnecessary alarms. These unnecessary alarms can be so frequent that they cause operational staff to ignore genuine alarms.

There are a number of third-party tools available for intrusion detection. For example, you can use Cisco's Intrusion Detection System (IDS) or ISS's RealSecure for real-time network traffic monitoring. IDSs are still in the process of being enhanced and developed.

Antivirus

The term "virus" is used to describe self-replicating computer programs that propagate themselves between files on a computer, and even between computers. Viruses usually, but not always, do something malicious, such as overwrite files or waste your bandwidth by sending copies of themselves to everyone in your address book.

Antivirus capabilities are a feature of some network and host-based firewalls. Network firewalls might inspect all incoming email traffic for virus-infected attachments, and filter them out. Host-based firewalls might change the configuration of the user's email client so that the email client sends all requests through the host-based firewall.

Firewalls are certainly not the only way to protect yourself from viruses, and if the firewall you choose doesn't have antivirus features, you'll need to complement it with antivirus software. The best way to protect your organization against viruses is to use a good-quality commercial antivirus package. These scanners examine the files, folders, mail messages, and Web pages on your computers, looking for the distinctive patterns of viral code. When the scanner detects something that looks like a virus, it quarantines the suspect object and warns you about what it found.

If your organization uses Outlook as its mail client, be sure to install the Outlook Security Update, which gives you a great deal of protection against email-borne viruses. (Note that this update's functionality is built into Outlook 2002, which comes as part of Office XP.)

TechNet has a great summary page that's continually updated to reflect newly emergent viruses. While you're reading it, reflect on the fact that most email-carried viruses spread because people do things they shouldn't, such as launch attachments from unknown sources.

VPNs and Encryption

Port forwarding is sufficient for publishing a Web site through your firewall. However, it's not sufficient if you want to connect two Internet-connected networks that are both protected by firewalls. For this, you should use a Virtual Private Network (VPN). A VPN is the extension of a private network that encompasses encapsulated, encrypted, and authenticated links across shared or public networks. VPN connections can provide remote access and routed connections to private networks over the Internet. Accessing the corporate network requires administrators to enforce strong authentication to validate identity as well as provide strong encryption to prevent users from communicating data "in the clear."

VPNs aren't strictly a firewall feature, and many businesses implement them by using completely separate, dedicated VPN devices. However, network architects generally place network firewalls at the perimeter of the network, just like a VPN. Both firewalls and VPNs are designed to improve network security, so it's logical that VPN capabilities have become a feature of many firewalls.

If you're using a Windows 2000 Server or Windows Server 2003 system as your network firewall, you already have VPN capabilities built into the base platform. All recent Windows platforms provide the authentication and encryption infrastructure to enable secure connectivity. With the Windows 2000 Server and Windows Server 2003 built-in VPN server and Windows XP VPN client, organizations can take advantage of a secure standards-based VPN directly "out of the box." Because Microsoft supports VPN standards such as L2TP/IPSec and smart card authentication, organizations have access to the encryption, authentication, and interoperability that best meet their VPN security needs. Although organizations often use VPNs to encrypt traffic over the Internet between users and the corporate network, they can also implement encryption between any Windows 2000, Windows Server 2003, and Windows XP machine. Since Microsoft has full standards-based support for the IPSec security extensions, organizations can provide robust encryption of all network traffic, without requiring cumbersome changes to deployed applications, servers, or network hardware.

Host-Based Firewalls

Host-based firewalls are software firewalls installed on each individual system. Depending on the software you choose, a host-based firewall can offer features beyond those of network firewalls, such as protecting your computer from spyware (a component of some free software that tracks your Web browsing habits) and Trojan horses (a program that claims to do one thing, but does another, malicious thing, such as recording your passwords). If you travel with a laptop, a host-based firewall is a necessity—you need protection wherever you connect to the Internet, and your hardware firewall can protect you only at home. Why would you buy third-party firewall software when Windows XP includes ICF for free? ICF is designed to provide basic intrusion prevention, but doesn't include the rich features of a third-party firewall application. Most third-party

firewalls protect you from software that could violate your privacy or allow an attacker to misuse your computer—features not found in ICF. Also, you can install third-party firewall programs on systems that have older versions of Windows. Note that firewall software doesn't replace antivirus software. You should use both.

Popular host-based firewall products include ZoneAlarm, Tiny Personal Firewall, Agnitum Outpost Firewall, Kerio Personal Firewall, and Internet Security Systems' BlackICE PC Protection. Most host-based firewall software is available in free or trial versions, so it won't cost you anything to download these packages and determine whether they meet your needs better than ICF.

Network Firewalls

Network firewalls protect an entire network by guarding the perimeter of that network. Network firewalls forward traffic to and from computers on an internal network, and filter that traffic based on the criteria the administrator has set. Network firewalls come in two flavors: hardware firewalls and software firewalls. Hardware-based network firewalls are generally cheaper than software-based network firewalls, and are the right choice for home users and many small businesses. Software-based network firewalls often have a larger feature set than hardware-based firewalls, and might fit the needs of larger organizations. Software-based firewalls can also run on the same server as other services, such as email and file sharing, allowing small organizations to make better use of existing servers. Network firewalls often include additional features that aren't necessary for host-based firewalls, as described in the following sections.

Proxy Services

If you have or are planning to have a home or small office network, you'll have to create a gateway from your firewall to the rest of the network. If you're implementing a software firewall on a specific computer, this means that you'll need at least two network cards in that machine. You attach one network card to the public interface (such as a DSL or cable modem), and You attach the other network card to your internal network. You then have to configure the computer to allow traffic on one side of the network to communicate with the other. ICS allows you to do this in both Windows 2000 and Windows XP.

However, at this stage in the game, many small office users decide to buy a dedicated residential gateway (see Figure 5). These units plug directly into the DSL router or cable modem and provide the functionality of a firewall and network hub. You need to configure a residential gateway to act in the stead of the computer running ICS when contacting the ISP. For example, if you had a static IP address, you would have to assign that IP address to the gateway instead of your computer. You could either assign a new IP address to your computer, or, more likely, instruct the computer to ask the gateway for an IP address.




Figure 5: A full-fledged small office network complete with a residential gateway

If a small business is using the 192.168.0.0 network ID for its intranet and its ISP has granted it the public address of $w1.x1.y1.z1$, then Network Address Translation (NAT) maps all private addresses on 192.168.0.0 to the IP address of $w1.x1.y1.z1$. If NAT maps multiple private addresses to a single public address, it uses dynamically chosen TCP and UDP ports to distinguish one intranet location from another.

Note: The use of $w1.x1.y1.z1$ and $w2.x2.y2.z2$ is intended to represent valid public IP addresses assigned by an ISP.

Figure 6 shows an example of using NAT to transparently connect an intranet to the Internet:

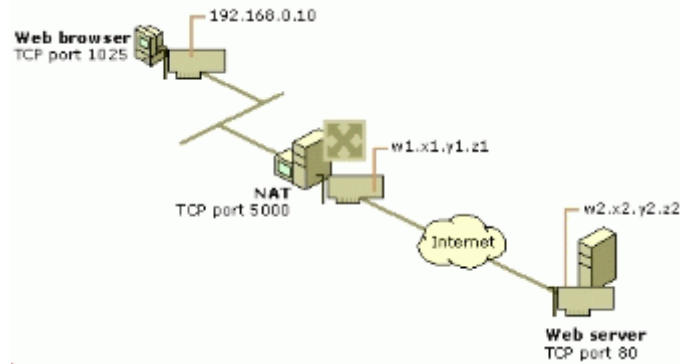


Figure 6: Using NAT to connect an intranet to the Internet

If a private user at 192.168.0.10 uses a Web browser to connect to the Web server at w2.x2.y2.z2, the user's computer creates an IP packet with the following information:

- Destination IP address: w2.x2.y2.z2
- Source IP address: 192.168.0.10
- Destination port: TCP port 80
- Source port: TCP port 5000

The private user's computer then forwards this packet to the NAT server, which translates the addresses of the outgoing packet to the following:

- Destination IP address: w2.x2.y2.z2
- Source IP address: w1.x1.y1.z1
- Destination port: TCP port 80
- Source port: TCP port 1025

The NAT server keeps the mapping of {192.168.0.10, TCP 1025} to {w1.x1.y1.z1, TCP 5000} in a table.

The NAT server then sends the translated packet over the Internet to the Web server. The Web server sends the response back to the NAT server. When the NAT server receives the packet, the packet contains the following public address information:

- Destination IP address: w1.x1.y1.z1
- Source IP address: w2.x2.y2.z2
- Destination port: TCP port 1025
- Source port: TCP port 80

The NAT server checks its translation table and maps the public addresses to private addresses and forwards the packet to the computer at 192.168.0.10. The forwarded packet contains the following address information:

- Destination IP address: 192.168.0.10
- Source IP address: w2.x2.y2.z2
- Destination port: TCP port 5000
- Source port: TCP port 80

For outgoing packets from the NAT server, the NAT server maps the source IP address (a private address) to the ISP allocated address (a public address), and maps the TCP/UDP port numbers to a different TCP/UDP port number.

For incoming packets to the NAT server, the NAT server maps the destination IP address (a public address) to the original intranet address (a private address), and maps the TCP/UDP port numbers back to their original TCP/UDP port numbers.

Note: NAT properly translates packets that contain the IP address only in the IP header. NAT might not properly translate packets that contain the IP address within the IP payload.

Reverse Proxy Services

Most proxy servers offer services beyond the standard functionality discussed above. Reverse proxy enables the firewall to provide secure access to an internal Web server (not exposing it to the outside) by redirecting external HTTP (application proxy) requests to a single designated machine. This isn't suitable for multiserver Web hosting (reverse hosting—described next—takes care of this), but it can be quite valuable when working with a single site.

Reverse hosting allows the firewall to redirect HTTP (application proxy) requests to multiple internal Web servers. One method/way is to provide access to multiple servers as subwebs of one large aggregate Web site or as multiple

independent Web servers. More flexible than reverse proxy but equally secure, this method enables you to abstract the physical architecture of your Web sites by mapping multiple servers to a single logical one. Both options allow the firewall to offer caching functionality, which can improve responsiveness. Server proxy provides the same functionality as reverse proxy and reverse hosting, but unlike these features, it works with protocols other than HTTP to provide secure access from the Internet to internal resources such as internal mail or SQL Server. To an outside user, the proxy server appears to be the mail or SQL Server. Basically, server proxy responds to external requests on behalf of the internal servers, which simply have to run the proxy client that redirects the listen directive on a given port to a proxy server. The security benefit is obvious: Placing servers behind a proxy prevents direct tampering from the outside and fools would-be attackers into thinking that the proxy server is the box containing the information they want.

Reverse proxy can be very useful. For instance, suppose you need to allow a Web server to query an internal database. There are several ways to do this. You could replicate the database to the outside (if it's not too large), but this puts the contents' integrity at risk. It might make more sense to move the Web and database servers behind the firewall and use reverse proxy or reverse hosting to get at the site. This option is very secure, although the overhead of running multiple Web servers behind the proxy might tax the proxy's ability to service Web requests from internal clients.

A third alternative is better yet: Place the Web server in the demilitarized zone (DMZ) and use the server proxy functionality of the firewall to query the database. This option, which Figure 7 (below) shows, provides good security and performance. Before you select any of these options, you should analyze your requirements so that you can balance necessary security against performance/usability.




Figure 7: Firewalls can act as reverse proxies for Web servers.

Firewalls for Small Offices and Home Offices

Firewalls used to be only for large corporate networks—but then again, Internet connections used to be only for large networks, too. Now that high-speed, always-on Internet connectivity is becoming more and more common, so too are attacks against connected computers. Firewalls help protect you against such attacks by screening out many types of malicious traffic. In addition, firewalls can help keep your computer from participating in attacks on others without your knowledge. The good news is that consumer-level firewalls provide good security without requiring that you be a computer security expert.

It used to be true that if you had a computer or two in a small office, the biggest risk you faced was losing data due to a fire, hardware failure, or other catastrophe. Although those risks are still with us, the blessing of always-on, high-speed Internet connectivity has exposed us to new threats, as well as intensifying some older ones. The good news is that, with the right tools, you can do a great deal to safeguard your computer systems against malicious attacks, viruses, and other bad stuff. Some of these tools come included with various versions of Windows. Others come from third-party vendors, such as Symantec, McAfee, and others. It's not necessarily important that you use a particular brand of tool; it's more important that you have the right tools, no matter who makes them.

Once you decide you need a firewall for your small office or home office, the first step in setting up a firewall is simple: Decide whether a hardware or a software solution will work best for your needs. Firewall products come in many different forms, from freely available software for your computer to tamper-resistant industrial units. Whether you buy a certified firewall or not, no matter what kind of

firewall you buy, all firewalls provide the same basic feature: control of inbound and outbound traffic.

When making this decision, you should be asking the following questions:

1. Are my computers all running Windows XP or Windows Server 2003? If so, you already have a built-in firewall, ICF, so you might not need to buy anything additional.
2. Do I want to share my Internet connection between multiple computers? If you do, you either have, or will have, your computers networked together. In that case, you can use Windows' Internet Connection Sharing (ICS) feature to share the connection. ICS is included with Windows 2000, Windows XP, and Windows Server 2003.
3. Do I want to be able to share my connection without using one computer as a firewall? Perhaps you want to share your Internet connection, but you don't want to have to use one particular computer as the gateway—ICS only works when the computer running it is powered up and on the network. Instead, you can buy an inexpensive appliance that acts as a gateway for sharing connections—these so-called “Internet access routers” or “residential gateways” almost always include basic firewall functionality.

For maximum security, you can install a hardware firewall to protect your small office network, and then combine it with ICF or another host-based firewall. You can configure host-based firewalls for individual machines so that each machine on your network can have different network permissions, if that's what you need. In addition, host-based firewalls can alert you when software on a machine that has a firewall is sending out data when it shouldn't. Hardware-based firewalls are typically very flexible and powerful; once you get them set up, you can leave them alone and let them work to silently protect you. Of course, getting them set up in the first place can be a difficult exercise!

Combining the two types can give you tighter security than using either one alone.

For maximum security, the most reliable way for small office users to protect a network is to purchase a router with firewall capabilities. These routers do more than act as a firewall—they network multiple computers, allow them to share a

single Internet connection, and might even support wireless networking. If you have more than one computer and an always-on broadband connection, a router/firewall gives you the benefits of a small office network and connects every computer to the Internet.

The router is generally a separate device from the cable or DSL modem—it's important to understand that most cable and DSL modems offer your small office network no protection whatsoever. If you didn't choose to pay extra for security features, you probably don't have any. If you're unsure about your modem, ask your ISP what level of protection your modem provides.

Using a router with firewall capabilities has several advantages over using host-based firewall software. Host-based firewalls can protect only one computer at a time, and configuring a host-based firewall for every computer on your small office network can be a nuisance. Host-based firewalls are completely unable to protect other types of devices connected to your network, such as a game system or personal television viewer. Finally, if you use multiple operating systems on your computer or experiment with beta operating systems, it's easy to forget to install a host-based firewall for every OS.

If you decide to use a hardware firewall, select one that has enough network ports to allow you to connect all computers and other network devices directly to it. As shown in Figure 8, wiring a firewall into your network is as simple as adding an answering machine to your phone line. Simply unplug the Ethernet connection between your cable/DSL modem and your PC, and plug it into the firewall. Then connect your computer and other network devices into your firewall. Configuring a hardware firewall isn't as simple—we'll discuss that later.




Figure 8: Wiring a firewall into a network

The following are some of the popular hardware firewall products available: Linksys Routers, NETGEAR Routers, and SMC Routers. Home and small office computers that are directly connected to the Internet require the added security of a firewall. The least expensive way to do this is to enable both ICF and ICS on a system, and allow all networked computers to connect through that system. You can enable ICS on only one Internet connection on your network, and you should protect this connection by enabling ICF. ICF can check only the communications that cross the Internet connection on which it's enabled. The following types of network topologies, which Figure 9 and Figure 10 show, are safe and the most recommended:




Figure 9: An example of using ICS to connect a LAN to the Internet.




Figure 10: An example of using ICS to connect multiple networks to the Internet.

You should avoid topologies with multiple Internet connections. If you must have multiple direct Internet connections on your network, you should ensure that ICF is enabled on each direct Internet connection in order to protect your network, as shown in Figure 11. However, because ICF works on a per-connection basis, this topology is still not a recommended topology because there's no central point of

administration through which you can ensure the continuous protection of all Internet connections.




Figure 11: Multiple Internet connections are more difficult to secure.

Likewise, as Figure 12 shows, providing Internet connectivity to your network by connecting your network hub directly to the Internet causes similar vulnerabilities and isn't a recommended topology.




Figure 12: LAN systems are not protected unless all traffic passes through ICF.

Enabling ICF on this type of network topology disrupts some network communications and provides protection only for the computer on which it's enabled. The other computers have direct connections to the Internet through the hub and aren't protected.

Firewalls for Enterprises

Organizations of all sizes want secure network connectivity to their business data and applications. The need to connect and collaborate with partners, customers, and remote/mobile employees anytime and anywhere has expanded network connectivity requirements beyond traditional wired local area networks (LANs) to include dial-up remote access, VPNs, and wireless networks. To enable greater access to the network and higher productivity, customers must address issues around security, management complexity, and cost. With Windows Server 2003, Windows 2000, Windows XP, and a carefully designed firewall architecture,

administrators can provide secure and integrated network connectivity to business-critical applications and data.

When addressing secure network connectivity, administrators need to consider the following:

- **Security:** Employees not only work from corporate offices, but also from branch offices, home offices, or the road. Providing remote connectivity requires solutions that are secure, standards-based, and manageable.
- **Management complexity:** Many vendors offer dedicated product solutions with little integration with other products and infrastructure. Setting up wireless clients with centralized authentication and policies can be a challenge unless there are integrated solutions.
- **Lowering cost:** Secure networking can be expensive if there are multiple products and technologies with separate licensing, support contracts, and training. For example, a secure VPN implementation might require a separate certificate authority for PKI, a separate authentication model, client-side software, and additional server gateways and firewalls.

By addressing these key secure connectivity challenges, organizations can achieve greater employee productivity, decrease costs, and improve business integration.

Host-Based Firewalls on Corporate Networks

Corporate networks also employ layers of defense. Often there will be some traffic screening at the router connecting the network to the Internet, one or more enterprise-class firewalls, virus scanning engines on the email servers, and some kind of intrusion detection mechanism.

An interesting consideration is whether host-based firewalls make sense in corporate networks. Enterprise firewalls and host-based firewalls operate at different defensive layers: Enterprise firewalls protect entire networks, whereas host-based firewalls protect individual hosts; so in one sense, combining the two seems appropriate. Host-based firewalls, however, block certain corporate network activity: For instance, some organizations periodically scan corporate clients for conformance to password policies; host-based firewalls interfere with this operation. Firewalls like ICF that block all incoming connections interfere with

LAN-based applications that need to send notices to client computers (printer status messages and Exchange new mail notification are two examples). Your organization's security policy needs to describe whether host-based firewalls are permitted and how they should be configured.

ICF in Windows XP obeys Network Location Awareness, and you can disable it via Windows Server 2003 Active Directory (AD) Group Policy. Mobile users don't have to remember to enable or disable ICF as they roam about. They can leave ICF enabled, being protected while at home or traveling. At the office, Group Policy can disable ICF whenever a computer is attached to the corporate network.

Using a Demilitarized Zone

Most organizations use their Internet connection to expose services to the public Internet. At a minimum, SMTP services are exposed to allow inbound email. You can use filtering and port forwarding to allow this traffic through a firewall, but many organizations require a DMZ to further protect the internal network. This section discusses how to secure your DMZ—the area in which you typically place your servers that expose public services to provide the best security.

A DMZ consists of front-end servers, back-end servers, and firewalls. The firewalls protect the front-end servers from the public network and filter traffic between the corporate network and back-end servers. A DMZ provides a multilayer protection system between the Internet and the internal network of an organization.

To provide protection, the DMZ comprises:

- A firewall that protects the front-end servers from Internet traffic.
- A set of “security-hardened” servers that support the services the application provides. You set up these servers so that dangerous Internet services, such as file sharing and Telnet, are disabled.
- A firewall that separates the back-end servers from the corporate networks and enables communication between the back-end servers and a few servers within the corporate network.

A DMZ is an important element for securing a site. You need to take additional security measures to protect data the back-end servers store. You can also store

extremely sensitive data or data that's needed elsewhere in your enterprise outside the DMZ, although doing so has negative performance implications and runs the risk, however small, of opening your corporate network to hacking. At the very least, a DMZ requires a router. A more sophisticated design would include two routers and a firewall. How complex your configuration needs to be depends on factors such as:

- How much security you need
- What sort of connectivity your system maintains to other networks (internal—corporate network; external—Internet)
- How many servers you need to protect

Figure 13 and Figure 14 below show DMZ configurations common in corporate environments today. Both provide excellent protection for the internal network. Figure 12 shows a simpler configuration; the configuration that Figure 14 shows protects the DMZ servers with the same security features used to protect the internal network, controlling DMZ access from the trusted and untrusted sides of the firewall. It's more complex but more secure.

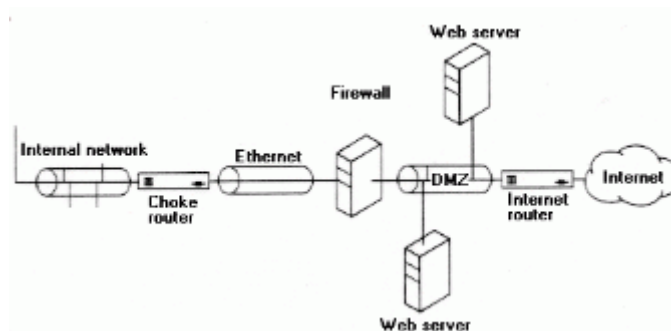


Figure 13: A simple DMZ configuration




Figure 14: DMZ configuration controlling DMZ access from the trusted and untrusted sides of the firewall

After DMZ topology, the most important step in securing the environment is controlling its traffic. You need to determine who's allowed to connect and who isn't, and then enforce those rules, usually with routers and firewalls. Routers can provide packet filtering, which controls traffic flow between two nodes, but this tends to decrease router performance, so you have to be careful not to overuse it. Check your router utilization before and after.

You must give particular attention to each server in the DMZ to ensure they're capable of withstanding malicious attacks. You can harden the exposed servers by using the Security Tools and Checklists for your servers' operating systems. You can also implement low-level filtering policies and close selective ports. For example, you should configure a host-based firewall on systems in a DMZ.

Standard DMZ Web Site Architectures

If you're going to implement an e-commerce or enterprise application, you have to be concerned with the security of your systems and data to ensure that people who shouldn't be accessing data can't get at it and to ensure that your system will be available despite attempts at a DoS attack. For enterprise applications, the main worry is unscrupulous employees—so security is typically enforced by using Windows and AD authentication and authorization.

But malicious attackers can also attack e-commerce applications from outside your company via the Internet. And since it's not practical to give every anonymous customer their own Windows logon ID, you'll need to use a different sort of authentication. Because the network is the Internet, instead of an intranet that you control, you'll also have to prepare your servers to make other sorts of attacks impossible or ineffective. Finally, you'll have to be especially careful protecting customer data, such as credit card numbers.

If you plan to host the site at your corporate facilities, you'll need to use a DMZ. The Internet-facing firewall must provide access to services such as HTTP, HTTPS, FTP, and SMTP mail. If you're collocating your servers at a hosting provider's network, a single Internet-facing firewall might be sufficient. However,

you'll also need to use a VPN to securely manage the site from your corporate network.

Figure 15 shows a typical architecture for hosting a large-scale e-commerce site using an enterprise's existing Internet connection.




Figure 15: A typical self-hosted e-commerce Web architecture

In general, here's what happens: Clients access the application over the Internet. Requests pass through a firewall, which filters out packets sent to the wrong address or wrong ports. The external firewall filters these requests, ensuring they originate from a valid address and are destined for a valid address and port number. If the firewall is an application-layer firewall, it will verify that the name of the page requested is valid and the request is well-formed. A Web server running IIS handles the requests, typically by using an ASP.NET page, and requests information from the database servers as needed. The Web servers may make requests to resources located within the corporate network. The DMZ's internal firewall—a final layer of protection for the internal network—filters these requests. This additional protection is critical since the risk of an attacker compromising externally facing Web servers is much higher than other internal servers, and an attacker might leverage them as a launching point for further attacks against the internal network.

Multilayer Firewall Web Site Architectures

Many organizations have security requirements that necessitate placing a firewall between the front-end Web servers and the back-end database servers. Figure 16 shows an example architecture that meets those requirements, and provides redundancy, while minimizing cost by using multihomed redundant firewalls. In this architecture, requests the Web servers send to the database servers must pass through the redundant firewalls. The firewalls can verify the source and

destination of the address, and validate that it's a legitimate request. This example architecture is placed at an Internet data center where administrators perform management of the systems remotely. Therefore, the firewalls have VPN capability, allowing administrators to securely access the Web and database servers from the corporate network.




Figure 16: Firewall placement in the Internet data center architecture

The firewalls in this architecture must support packet filtering, stateful inspection, application-level filtering, and VPNs. Further, they must be capable of supporting a sufficient number of connections and throughput to handle inbound requests from Web users, outbound responses to clients, requests between the Web and database servers, and connections across the VPN for system management. A typical production architecture would include another network connected to the firewalls or directly to the servers for backing up data.

Firewall Products

This section provides an overview of some key firewall products and categories.

Internet Connection Firewall

ICF is host-based firewall software that you use to set restrictions on what traffic is allowed to enter your computer. ICF protects your computer against external threats by allowing safe network traffic to pass through the firewall into your network, while denying the entrance of unsafe traffic.

If you configure ICF on a system that uses ICS, that protection is extended to all traffic passing through the system. If your network uses ICS to provide shared

Internet access to multiple computers, you should enable ICF on the shared Internet connection, as Figure 17 shows:

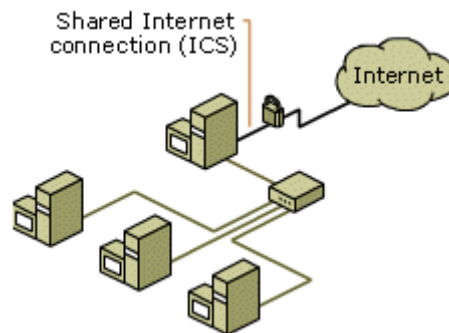


Figure 17: Enabling ICF on an ICS-enabled system increases the security of the entire network.

Although it's critical to enable ICF on the Internet connection of any computer that's connected directly to the Internet, you should enable ICF on all network interfaces to protect against attacks originating from the internal network. ICF can provide protection to a single computer connected to the Internet with a cable modem. To check to see whether ICF is enabled on Windows Server 2003, or to enable ICF, see [To enable or disable ICF](#). For information about using ICF, see [Protecting your home or small office network using Internet Connection Firewall](#).

How ICF Works

ICF is considered a stateful firewall. To prevent unsolicited traffic from the public side of the connection from entering the private side, ICF keeps a table of all communications that have originated from the ICF computer. When used in conjunction with ICS, ICF tracks all traffic that has originated from the ICF/ICS computer and all traffic that has originated from private network computers. ICF compares all inbound traffic from the Internet against entries in the table. ICF allows inbound Internet traffic to reach the computers in your network only when there's a matching entry in the table that shows that the communication exchange originated from your computer or private network.

To thwart common hacking attempts (such as port scanning), the firewall drops communications that originate from the Internet. Rather than sending you notifications about firewall activity, ICF silently discards unsolicited communications, because frequently sending such notifications could become a

distraction. Instead, ICF creates a security log to track this activity. For more information, see Internet Connection Firewall security log.

You can configure services to allow the ICF computer to forward unsolicited traffic from the Internet to the private network. For example, if you're hosting an HTTP Web server service, and the HTTP service is enabled on your ICF computer, the ICF computer forwards unsolicited HTTP traffic. ICF requires a set of operational information, known as a service definition, to allow it to forward unsolicited Internet traffic to the Web server on your private network. Service definitions for ICF work on a per-connection basis. If your network has multiple firewall connections, you should configure service definitions on all firewall connections. For information about service definitions, see Service definitions for Internet Connection Firewall and Internet Connection Sharing. For information about adding and editing service definitions, see To manage service definitions for ICF or ICS.

Note: Microsoft includes ICS and ICF only with Windows XP; Windows Server 2003, Standard Edition; and the 32-bit version of Windows Server 2003, Enterprise Edition. These features aren't included with Windows Server 2003, Web Edition; the 32-bit version of Windows Server 2003, Datacenter Edition; or the 64-bit versions of the Windows Server 2003 family.

For those who live and breathe TCP/IP, ICF's engine uses addresses, ports, sequence numbers, and flags in its state table. For TCP, the outgoing request must have only the SYN flag turned on; incoming replies must have only the ACK and SYN flags turned on; the next outgoing packet must have only the ACK flag turned on. If this sequence is violated, ICF terminates the connection. Also, ICF drops any incoming packet (including ACK-SYN) that it can't associate with (using addresses, ports, and sequence numbers) a previous outgoing ACK. ICF also drops any incoming unsolicited SYN unless it matches a user-defined exception (see the next section). When ICF discards packets, it does so silently; it never returns an RST.

ICF "forgets" about a particular state between the client and a server when:

- It sees a special connection-termination sequence for TCP-based communications (ACK-FIN, ACK, ACK-FIN, and ACK)

- A period of inactivity (a “timeout”) for UDP-based communications

Essentially, ICF allows in only that which is a reply to a previous request that went out. ICF blocks and discards all other incoming traffic. It seems simple, but it’s an extremely effective method for protecting a computer. Even enterprise-class firewalls operate using the same basic principle. This method means that ICF blocks the following kinds (among others) of potentially dangerous communications:

- **Scans.** Attackers often scan computers looking for vulnerabilities, especially the popular well-known cable-modem subnets. Because incoming scans are “unsolicited” (i.e., don’t match something in ICF’s traffic memory), ICF blocks them.
- **Many (but not all) Trojans.** Say you get infected with a Trojan horse program. Many of these announce their existence to some database somewhere. If an attacker tries to connect to the Trojan on your computer, ICF blocks it. Note that this applies only to Trojans in which the attacker makes the first connection to the infected computer; other Trojans that make the first connection to the attacker open a connection in ICF’s memory, allowing the attacker to reply. This is why you need a virus scanner in addition to ICF—good-quality virus scanners also prevent you from getting infected with Trojans in the first place.
- **File sharing and anonymous connections.** Windows networking is intended to allow easy file sharing between computers; you use anonymous connections for discovering a computer’s name and list of available file shares. Of course, on the Internet you really don’t want to do this; ICF prohibits these kinds of connections.

By default, turning on ICF automatically blocks all communications originating from foreign computers to all ports on your computer. This essentially renders your computer invisible to port scanners, hacking tools that repetitively try different ports on a network address to see how a particular machine might be attacked. Hackers frequently run port scanners against DSL and cable-modem services to see whether they can find any unsecured machines; with ICF running, the port scanners never see your machine.

If you want to enable Internet users to communicate with your machine (as you might if you’re running a Web server or using a popular peer-to-peer program),

you can choose to allow communications to specified ports by clicking on the Settings button on the Advanced tab in the properties window of your public network connection. Although the Code Red and Nimda viruses demonstrate the importance of patching services even when they're not exposed through the firewall, once you open up a port to the outside world, it's imperative you apply and keep up-to-date with any patches for that service. Windows Update provides an easy way to ensure that you're up-to-date with Windows XP services. If you're hosting a service that you want other computers to have access to, you need to let ICF know that it's OK for computers to access it. The classic example is a Web server. ICF is easily configurable for any services that you might be running on your computer and comes preconfigured with definitions for common protocols and services such as HTTP, FTP, and SMTP. Even if the service or programs that you want to expose aren't on the default list, you can add them by defining the internal and external TCP ports that they utilize.

ICF and Notification Messages

Because ICF inspects all incoming communications, some programs, especially email programs, might behave differently when ICF is enabled. Some email programs periodically poll their email server for new mail, and some email programs wait for notification from the email server.

Outlook Express, for example, automatically checks for new email when its timer tells it to do so. When new email is present, Outlook Express prompts the user with a new email notification. ICF won't affect the behavior of this program because the request for new email notification originates from inside the firewall. The firewall makes an entry in a table noting the outbound communication. When the mail server acknowledges the new email response, the firewall finds an associated entry in the table and allows the communication to pass, and then the user receives notification that a new email message has arrived.

Outlook 2000, however, is connected to a Microsoft Exchange server that uses a remote procedure call (RPC) to send new email notifications to clients. Outlook 2000 doesn't automatically check for new email when it's connected to an Exchange server. The Exchange server notifies Outlook 2000 when new email arrives. Because the Exchange server—which is outside the firewall—not

Outlook 2000—which is inside the firewall— initiates the RPC notification, ICF can't find the corresponding entry in the table, and it doesn't allow the RPC messages to cross from the Internet into the network. ICF drops the RPC notification message. Users can send and receive email, but need to manually check for new email.

Advanced ICF Settings

The ICF security-logging feature provides a way to create a security log of firewall activity. ICF is capable of logging both traffic that's permitted and traffic that's rejected. For example, the firewall, by default, doesn't allow incoming echo requests from the Internet. If the Internet Control Message Protocol (ICMP) **Allow incoming echo request** isn't enabled, then the inbound request fails, and ICF generates a log entry that notes the failed inbound attempt. For information about ICMP, see Internet Control Message Protocol (ICMP). ICMP allows you to modify the behavior of the firewall by enabling various ICMP options, such as **Allow incoming echo request**, **Allow incoming timestamp request**, **Allow incoming router request**, and **Allow redirect**. The ICMP tab provides brief descriptions of these options. For navigation and instructions for ICMP, see Enable Internet Control Message Protocols.

You can set the allowable size of the security log to prevent the potential overflow that DoS attacks could cause. ICF generates event-log entries into the Extended Log File Format as established by the World Wide Web Consortium (W3C). For more information about ICF security logging, see Internet Connection Firewall security log file overview.

Other Host-Based Firewalls

If you're not running Windows XP, or if you want to have greater control (and awareness) of what your firewall is doing on your behalf, a separate host-based firewall software package might serve you better. There are a variety of good products available that enhance your computer's security. For example, ZoneAlarm by Zone Labs not only filters incoming connections, but also filters outgoing connections by program. That means that you can specify which programs on your computer should be able to communicate over the Internet and which, if any, should be prevented from doing so.

Microsoft Internet Security and Acceleration Server 2000

ISA Server is part of the Microsoft .NET Enterprise Server family, which comprises a comprehensive set of server applications for quickly building, deploying, and managing scalable and integrated Web-based solutions and services. Designed with mission-critical performance and integration in mind, the .NET Enterprise Servers are built from the ground up for interoperability using open Web standards such as XML. The .NET Enterprise Servers, along with the Windows 2000 operating system, supply the foundation for the .NET platform, which enables the third-generation Internet: where software is delivered as a service; is accessible by any device, at any time and any place; and is fully programmable and customizable. Microsoft explicitly designed the .NET platform to enable the rapid development, integration, and orchestration of any group of Web services and applications into a single comprehensive solution.

ISA Server is an extensible enterprise firewall and Web-cache server that integrates with Windows 2000 for policy-based security, acceleration, and management of Internetworking. ISA Server provides two tightly integrated modes: a multilayer firewall and a high-performance Web-cache server. The firewall provides filtering at the packet, circuit, and application layers; stateful inspection to examine data crossing the firewall; control of access policy; and routing of traffic. The cache improves network performance and the user experience by storing frequently requested Web content. You can deploy the firewall and cache on dedicated servers separately, or integrated on the same box. Sophisticated management tools simplify policy definition, traffic routing, server publishing, and monitoring. ISA Server builds on Windows 2000 security, directory, VPN, and bandwidth control. Whether deployed as a set of separate firewall and cache servers or in integrated mode, ISA Server can enhance network security, enforce consistent Internet usage policy, accelerate Internet access, and maximize employee productivity for organizations of all sizes.

ISA Server 2000 Enterprise Edition is Microsoft's scalable enterprise firewall and Web-caching server. Microsoft designed ISA Server Enterprise Edition to meet the performance, management, and scalability needs of high-volume Internet traffic environments with centralized server management, multiple levels of

access policy, and fault tolerance. ISA Server Enterprise Edition provides fast, secure, and scalable Internet connectivity for mission-critical environments. Administrators must secure the network access points of corporate networks against hackers and unauthorized access. Blocking traffic and shutting down ports aren't sufficient or feasible in an Internet-connected organization. Having security solutions that "look inside" network traffic to validate application-specific requests mitigates risks. ISA Server Enterprise Edition provides organizations with the stateful-packet inspection and application-layer firewall protection required to protect against today's sophisticated attacks. With ISA Server Enterprise Edition's application-level filtering technology, you can mitigate attacks such as Code Red and Nimda at the firewall before they enter company networks.

ISA Server Enterprise Edition integrates with Microsoft Management Console (MMC) and AD to provide a single directory to validate and manage all access requests for application data or services. This enables consolidation of access control and authorization policy in a centrally managed, replicated, and secure repository. ISA Server Enterprise Edition is also designed to work best with Exchange 2000 and IIS to provide fast and secure access to email and Web content.

The following sections describe some of the more significant product capabilities of ISA Server.

Multilayer Firewall Security

A firewall can enhance security through various methods, including packet filtering, circuit-level filtering, and application filtering. Advanced enterprise firewalls, such as that provided with ISA Server Enterprise Edition, combine all three of these methods to provide protection at multiple network layers.

Circuit-Level Filtering

At the circuit level, the ISA Server Firewall service works with virtually all Internet applications and protocols—such as Telnet, mail, news, Microsoft Windows Media technologies, RealAudio, and Internet Relay Chat (IRC)—and other client applications. The Firewall service makes these applications perform as if they

were connected directly to the Internet. ISA Server offers circuit-level filtering for both firewall and SecureNAT clients.

Circuit-level filtering enables support for virtually all standard and custom Internet applications on the Windows platform. These applications communicate on the network by using Winsock, and you can support them, unmodified, on client machines that have the Firewall client software installed.

Circuit-level filtering inspects sessions, rather than connections or packets. A session can include multiple connections, providing a number of important benefits for Windows-based clients running the Firewall client software.

Packet Filtering

The packet-filtering capability of ISA Server enables the administrator to control the flow of IP packets to and from ISA Server. When packet filtering is enabled, ISA Server drops all packets on the external interface unless they're explicitly allowed, statically, by IP packet filters, or dynamically, by access policy or publishing rules.

IP packet filtering intercepts and evaluates packets before they're passed to higher levels in the firewall engine or to an application filter. You can configure IP packet filters so that only specified packets will be passed through ISA Server. This practice provides a high level of security for the network. IP packet filtering can block packets originating in specific Internet hosts and can reject packets associated with many common attacks. IP packet filtering can also block packets destined to any service on an internal network, including the Web proxy, a Web server, an SMTP server, and others.

IP packets filters are static—communication through a given port is always either allowed or blocked. Allow-filters allow the traffic through, unconditionally, at the specified port. Block-filters always prevent the packets from passing through the ISA Server computer.

ISA Server supports dynamic packet filtering, opening ports automatically only as required for communications, and closing the ports when the communication ends. This approach minimizes the number of exposed ports in either direction and provides a high level of security for a network.

ISA Server supports inbound and outbound IP packet filtering. ISA Server's packet filtering also allows for blocking fragments and detecting packet-level attacks against the firewall.

Application-Level Filtering

The most sophisticated level of traffic inspection that the ISA Server firewall provides is the application-level security. "Smart" application filters can analyze a data stream for a given application and provide application-specific processing, including inspecting, screening or blocking, redirecting, or even modifying the data as it passes through the firewall. This mechanism protects against known exploits such as unsafe SMTP commands or attacks against internal Domain Name System (DNS) servers. Third-party tools for content screening, including virus detection, lexical analysis, and site categorization, also use application and Web filters to further extend the firewall.

Stateful Inspection

Stateful inspection examines data crossing the firewall in the context of its protocol and the state of the connection. At the packet level, ISA Server inspects the source and destination of the traffic indicated in the IP header and the port in the TCP or UDP header identifying the network service or application used. Dynamic packet filters enable the opening of a port only in response to a user's request and only for the duration required to satisfy that request, reducing the vulnerability associated with open ports. ISA Server can determine dynamically which packets can be passed through to the internal network's circuit- and application-layer services. Administrators can configure access-policy rules that open ports automatically only as allowed and then close the ports when the communication ends. This process, known as dynamic packet filtering, minimizes the number of exposed ports in either direction and provides a high level of problem-free security for the network.

Integrated Intrusion Detection

With the help of technology that a firm known as Internet Security Systems provides, ISA Server can help administrators identify and respond to common network attacks such as port scanning, WinNuke, and Ping of Death. This

technology provides ISA Server with an integrated intrusion-detection mechanism that identifies these kinds of attacks. The alert also specifies what action ISA Server should take when an attack is recognized, action that might include sending an email message or a page to the systems administrator, stopping the Firewall service, writing to the system event log, or running any program or script. With additional help from third parties, ISA Server can help administrators identify and respond to other common network attacks.

High-Performance Web Cache

ISA Server has a completely redesigned Web cache that enables it to place cache into RAM. This high-performance Web cache provides greater scalability on the back end as well as a faster overall Web-client response time.

Cache Array Routing Protocol

ISA Server uses the Cache Array Routing Protocol (CARP) to provide seamless scaling and high efficiency in an array of multiple ISA Server computers. CARP uses hash-based routing to provide a deterministic “request resolution path” through an array. Having a request resolution path, which is based on a hashing of array member identities and URLs, means that for any given URL request, the browser or downstream proxy server can know exactly where the requested information is stored in the array.

Chained-Configuration Cache Placement

In this context, the term “chaining” refers to a hierarchical connection between individual ISA Server computers or arrays of ISA Server computers. With chaining, an ISA Server sends client requests upstream through the chain of cache servers until the requested object is found. During this process, each ISA Server caches the object as it is forwarded through each server to the client. Consequently, chaining becomes an effective means of distributing server load and fault tolerance.

You can use the chained configuration to position content closer to users who need it, resulting in faster Web-client response times and reduced Wide Area Network (WAN) traffic. ISA Server makes a distributed Web cache possible in

which users can obtain Web pages from ISA Server rather than from individual Web sites.

Active Caching

With a feature known as active caching, whereby you can configure ISA Server to automatically update objects in a cache, ISA Server can optimize bandwidth utilization by proactively refreshing content. With active caching, ISA Server proactively updates objects that a client accesses frequently during periods of low network traffic.

Active caching is a way to keep objects fresh in the cache by verifying them with the originating Web server before the objects expire and a client accesses them. The goal is to expedite those client accesses that would ordinarily require a round-trip to the originating server to revalidate the data. Because there's a cost associated with this (in both proxy processing and network bandwidth), the goal is to refresh only those objects that a client is likely to access in the future. In contrast, object "popularity" isn't a useful criterion for this because many popular pages never expire. This is due to the fact that clients refresh the pages manually to keep the data fresh. In addition, an object might be popular for only a short time. The active-caching code tries to identify objects that follow precisely the pattern of accessed content that active refreshing would help—that is, objects that expire and which a client then touches again.

Unified Management

ISA Server takes advantage of the Active Directory service, VPN, and MMC. All these capabilities, especially MMC, help to make administration easier because operations personnel are familiar with them and can manage both the firewall and Web cache from one console.

Using built-in reporting tools, ISA Server supports the running of scheduled standard reports that detail Web usage, application usage, network-traffic patterns, and security. ISA Server provides extensive support for reporting such matters as frequency of Internet access, what is being accessed, and by whom. By alerting and reporting such matters as out-of-boundary activity to administrators, ISA Server can help them to better understand how employees

are using the Web. This is helpful for capacity planning and enforcing corporate policies.

For example, with modification to the AD schema and creation of policy, administrators can configure ISA Server to run in an array. Arrays allow administrators to apply policy, such as port configuration, at the enterprise level, thereby allowing them to apply a single change to every ISA Server within each array. AD multimaster replication of ISA Server configuration ensures that each server in an array receives current and automatically updated configurations. Administrators must control and enforce security policies while simultaneously supporting employees who need Internet access. Thanks to its integration with AD, ISA Server provides comprehensive support for controlling access by user, group, application, destination, content type, and schedule.

Enterprise Policy and Access Control

ISA Server also supports the creation of enterprise-level and local array policies, for centralized or local enforcement. You can install ISA Server as a stand-alone server or as an array member. For easier management and administration, array members share the same configuration. When you modify the array configuration, you also modify all the ISA Server computers in the array, including all their access and cache policies.

Centralized administration can also mean greater security. Administrators can configure a single server, and apply that configuration to multiple ISA Servers. An enterprise can take this centralized management one step further, allowing administrators to implement one or more enterprise policies, which include site and content rules and protocol rules. Administrators can apply an enterprise policy to any array and augment it with the array's own policy. This approach enforces enterprise policies at branch and departmental levels while allowing local administrators to further restrict access.

Nortel Networks Alteon Switched Firewall

One of Microsoft's internal networks uses Alteon hardware firewalls, and serves as an interesting study. It contains a high-speed security network designed to support enterprise-scale applications such as online transaction processing

(OLTP) and business intelligence (BI). Nortel Networks Alteon Switched Firewalls (ASFs) that are configured for high availability provide internal network security. Two separate ASFs provide stateful failover and deep packet inspection of all connections passing through them. The ASFs provide 3.2 gigabytes of data throughput and 32,000 sessions per second with 500,000 concurrent sessions. Microsoft's configuration consists of two Alteon Firewall Accelerators and two Alteon Firewall Directors. Environments with higher capacity requirements can expand this configuration up to twelve Alteon Firewall Directors; the ASF Plug and Play (PnP) configuration features provide the ability to scale network security with zero downtime.

From an operational perspective, the ASF runs Check Point Firewall-1 Next Generation software much the same as any server running Check Point software. Procedures for monitoring security events, handling security management, and so on are the same.

What's different is the scale and throughput of the firewalls. The increased scale introduces some operational differences in how you manage the system, including features unique to the characteristics of the ASF and some features in the Check Point Firewall-1 Next Generation software that are not in other versions of Check Point Firewall software. These unique differences are outlined below.

Firewall Scale Considerations

The Nortel Networks ASF has some unique capabilities, notably DMZ partitioning and PnP scaling, that allow for both scale-up and scale-out of the firewall solution. Nortel Networks provides training on such capabilities as part of the system installation process. Details on the operational considerations of managing ASFs are provided later in this chapter.

Large networks are typically very robust in design, and effective management can be difficult. Nortel Networks provides a single management solution for the ASFs by using the Check Point software management suite.

When operating a large network, the firewalls must scale to take advantage of the capacity of the server and application requirements. The firewalls must be able to scale in bandwidth, and you must maintain a very secure environment for

the traffic within the data center while still providing access to external resources. This does impose some requirements on the security administrator; for example, the firewalls need to support and scale the number of add/drop sessions per second and concurrent number of sessions while maintaining adequate throughput.

When running multiple firewalls, whether multiple instances of the same firewall or a number of different firewalls working together, the key requirement is that the firewalls not interfere with one another and maintain state between themselves for failover and failback purposes. Nortel Networks provides a single image configuration for the Alteon Firewall Cluster that provides the security administrator with the ability to perform upgrades and other administrative tasks on a single firewall, while synchronization procedures populate the changes to the rest of the cluster.

Alteon Switched Firewall Considerations

The Alteon Switched Firewall contains a number of features, such as multilink trunking and VLAN tagging, that affect certain design aspects of the network. Nortel Networks provides installation support options that you can negotiate with a local Nortel Networks sales representative. For more information, visit <http://www.nortelnetworks.com/index.html>.

Tools

Check Point Enterprise Management Client is the suite of management tools that you need to manage the ASFs. Within the Check Point Management software, there are three main modules that deal with different operational aspects of the ASF:

- Check Point Policy Editor
- Check Point Log Viewer
- Check Point System Status Viewer

These modules are described in the following sections.

Check Point Policy Editor

The Check Point Policy Editor provides comprehensive management of the Check Point Firewall engine that runs on the ASFs. The Policy Editor enables the security administrator to perform a wide range of functions, including:

- Create, delete, and edit security policies on the ASF cluster
- Create, delete, define, and edit objects, protocols, servers, and users within the firewalls
- Install and uninstall the security policies assigned to the ASF clusters
- Create, delete, and edit address translation rules from the ASF clusters

Check Point Log Viewer

The Log Viewer provides administrators a visual representation of the traffic flow through the ASFs. Security policies define which traffic to add to the Log Viewer, so communications that don't specify logging aren't logged. The Log Viewer is capable of displaying the following information for each communication:

- Date and time
- Interface on which the traffic entered
- The action taken on the communication
- The rule number
- Encryption and address translation details

The Log Viewer has the ability to filter on each of the information categories above to focus on any element of the logs. It has three modes: log, account, and active. The log mode shows each connection as it's created. The account mode shows information pertinent to accounting, such as the duration of the connection and the number of bytes transferred. The active mode lists all the connections currently open through the ASF. When the active mode is enabled, the administrator can block any connection (which can be very useful when policing the security domain).

Check Point System Status Viewer

The System Status Viewer gives a graphical view of the working state of each firewall, and provides information on how many packets have been accepted,

logged, and dropped. You can also use this tool to summarize the policy for a particular firewall. In addition, the System Status Viewer is the system monitoring interface that allows the generation of alerts under security administrator–defined conditions.

Distributed Enterprise Management

Distributed Enterprise Management allows the ASF administrator to configure and support multiple firewalls from one workstation. The configuration of other activities is done through their respective GUIs and saved to a central management server. The changes or actions applied to the required firewall are then downloaded from the management server, which allows for scalable management as the number of firewalls under the administrator’s control increases. The management server holds a database that contains all the information needed to generate the required security policies.\

Object Creation and Management

After the installation of the ASFs and after the Check Point Management software is complete, the next step is to create and apply the rule-base of the security policy. The security of the ASF is centered on an object-oriented rule-base, which you manage through the Policy Editor’s management menu and tool bars. The utilities required to create and manage the different objects are in the following tools:

- Network Objects Manager
- Services Manager
- Resources Manager
- Servers Manager
- Users Manager
- Time Objects Manager

Depending on the size of a network, including subnets used, range of services provided, and number of resources and protocols, you can save a significant amount of time by defining each element in the Objects database. When building the Objects database, you should adopt standards for object creation. Some standards to consider are the following:

- Naming convention
- Color
- Groups

Assigning a short, intuitive name for an object can be a difficult task. Therefore, when truncating a name to a maximum of 10 characters or symbols, you should adopt a naming convention. For example:

WS31OWSEWA (WS=Web Server, 31=Thirty-one, OW=Outlook Web Access, SE=Seattle, WA=Washington)

To many people the name WS31OWSEWA might seem meaningless, but to the security administrators it would make sense.

The ability to assign a color to an individual object is a powerful administrative feature. Color can show certain characteristics of an object, such as the department it belongs to, or where it's located. Consider the following color scheme:

- Red = Firewalls
- Green = Internal protected objects
- Yellow = External possibly harmful objects
- Black = Services

Operating

The following sections describe some of the software-specific operations tasks that you can perform on the ASF. Detailed documentation on how to perform these tasks is provided with the ASF and ASF training, and is not replicated here.

Configuration

You can configure the ASF as one or more "ASF offerings." An ASF offering is simply a set of hardware components you use together to create the ASF.

You can run the entire ASF as a single firewall, or you can split the system into multiple DMZ partitions, each with its own copy of the set of rules and policies.

The Nortel Networks installation contract, documented and provided to the customer, covers all hardware configurations. You can configure the Policies, Objects, and Rules bases by using the Check Point Management Client software suite of tools.

Maintenance

The only periodic maintenance required on the ASFs is replacement of the licenses for the Check Point firewall engines on the management station every 12 months, depending on environmental conditions within the data center. In the event of any hardware problems, Nortel Networks provides maintenance under the terms of the applicable service and support agreement.

Real-Time Monitoring

The Check Point Management software suite provides security-log monitoring. This feature provides policy-based definition and event notification to monitor the security settings of the ASFs and Check Point Firewall-1 software environment. Monitoring is continuous and in real-time to provide rapid notification of fault conditions and to ensure that all critical rules and policies are operating within defined limits.

System health monitoring provides a real-time view of the ASF, which:

- Allows the security administrator to check policy configuration and status information via a graphical user interface.
- Offers a consolidated view of the rules and objects in the ASFs.
- Provides management wizards, which enable an administrator to implement system changes across all jASFs without having to make modifications multiple times on multiple firewalls.
- Provides a Security Logs status page, from which a security administrator can view a 24-hour history of all the logs the ASF has generated.

Supporting

Incidents that affect a limited number of users might be indicative of larger problems that can affect the overall functionality of a network. Troubleshooting is an important part of incident prevention and problem management.

Troubleshooting guidance is provided to help operations personnel understand various incidents and problems that might occur.

Nortel Networks provides support of the ASF hardware under the terms of a service agreement. There are various options, depending on the needs of the data center.

Capacity Management

Whether you configure the firewall as one large DMZ or multiple DMZs, capacity management is important. Being an ASIC-based switch, the ASF can perform wire speed routing and intelligent decision making to provide the best performance for a given network.

Nortel Networks offers extensive performance analysis and capacity planning services to ensure clients obtain the optimum throughput from ASFs. For more information, see <http://www.nortelnetworks.com/>.

There are several other vendors for hardware firewalls that deserve investigating. Among these are SonicWALL, Nokia, Cisco, WatchGuard, NetScreen Technologies, Lucent, and Nortel, which all offer one or more models of hardware firewalls.

Summary

Of all the choices presented here, the only one you can really go wrong with is choosing none of them. Don't make the mistake of thinking that no one will attack your network, because with the rise in automated attack tools, your network is as much at risk as every other network on the Internet. Unfortunately, there's a strong subculture that misaligns the bravado of a deep understanding of how networks and computer security works with digital breaking and entering. Because of the popularity and ease-of-use of always-on connections, small office networks provide easy targets. However, by taking some simple precautions, you can protect network and your data.

To take advantage of the networked world, organizations must prevent unauthorized users from accessing their networks, and at the same time, ensure that authorized users have access only to authorized assets. By providing advanced security technologies, common management, and lower cost through integrated solutions, Microsoft can enable businesses to take advantage of the network connectivity.

Security is one of the most important features of a business site. You can implement security at all levels: network, platform, application, and database. You need to implement security technologies and tools that are cost-effective and safe, and that don't adversely affect the performance of the site.

You must implement a DMZ in your site, and configure firewalls that prevent access to dangerous services. Network segregation, data encryption, and intrusion detection will provide further security for your network.